


07-05-05

ATF ZZW

DOCKET NO.: MSFT-0160 / 142385.01

PATENT



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

**In Re Application of:**

Luis Felipe Cabrera, Deborah C. Jones,  
Ravisankar Pudipeddi and Stefan R. Steiner

**Confirmation No.:** 5398

**Application No.:** 09/644,667

**Group Art Unit:** 2167

**Filing Date:** August 24, 2000

**Examiner:** Mohammad Ali

**For:** PARTIAL MIGRATION OF AN OBJECT TO ANOTHER STORAGE LOCATION  
IN A COMPUTER SYSTEM

**EXPRESS MAIL LABEL NO:** EV670669028US

**DATE OF DEPOSIT:** July 1, 2005

EV670669028US

MS Appeal Brief - Patent  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

**APPEAL BRIEF TRANSMITTAL  
PURSUANT TO 37 CFR § 1.192**

Transmitted herewith in triplicate is the APPEAL BRIEF in this application with respect to the Notice of Appeal received by The United States Patent and Trademark Office on **May 2, 2005**.

- ☐ Applicant(s) has previously claimed small entity status under 37 CFR § 1.27 .
- ☐ Applicant(s) by its/their undersigned attorney, claims small entity status under 37 CFR § 1.27 as:
- ☐ an Independent Inventor
  - ☐ a Small Business Concern
  - ☐ a Nonprofit Organization.
- ☐ Petition is hereby made under 37 CFR § 1.136(a) (fees: 37 CFR § 1.17(a)(1)-(4) to extend the time for response to the Office Action of \_\_\_\_\_ to and through comprising an extension of the shortened statutory period of \_\_\_\_\_ month(s).

	SMALL ENTITY		NOT SMALL ENTITY	
	RATE	FEE	RATE	FEE
<input checked="" type="checkbox"/> APPEAL BRIEF FEE	\$250	\$	\$500	\$500.00
<input type="checkbox"/> ONE MONTH EXTENSION OF TIME	\$60	\$	\$120	\$
<input type="checkbox"/> TWO MONTH EXTENSION OF TIME	\$225	\$	\$450	\$
<input type="checkbox"/> THREE MONTH EXTENSION OF TIME	\$510	\$	\$1020	\$
<input type="checkbox"/> FOUR MONTH EXTENSION OF TIME	\$795	\$	\$1590	\$
<input type="checkbox"/> FIVE MONTH EXTENSION OF TIME	\$1080	\$	\$2160	\$
<input type="checkbox"/> LESS ANY EXTENSION FEE ALREADY PAID	minus	(\$ )	minus	(\$ )
TOTAL FEE DUE		\$0		\$500.00

- ☒ The Commissioner is hereby requested to grant an extension of time for the appropriate length of time, should one be necessary, in connection with this filing or any future filing submitted to the U.S. Patent and Trademark Office in the above-identified application during the pendency of this application. The Commissioner is further authorized to charge any fees related to any such extension of time to Deposit Account 23-3050. This sheet is provided in duplicate.
- ☐ A check in the amount of \$     .00 is attached. Please charge any deficiency or credit any overpayment to Deposit Account No. 23-3050.
- ☒ Please charge Deposit Account No. 23-3050 in the amount of \$500.00. This sheet is attached in duplicate.
- ☒ The Commissioner is hereby authorized to charge any deficiency or credit any overpayment of the fees associated with this communication to Deposit Account No. 23-3050.

Date: July 1, 2005



Jeremiah J. Baunach  
Registration No. 44,527

Woodcock Washburn LLP  
One Liberty Place - 46th Floor  
Philadelphia PA 19103  
Telephone: (215) 568-3100  
Facsimile: (215) 568-3439

DOCKET NO.: MSFT-0160 / 142385.1

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of:

**Luis Felipe Cabrera, Deborah C. Jones,  
Ravisankar Pudipeddi, Stefan R. Steiner**

Confirmation No.: **5398**

Application No.: **09/644,667**

Group Art Unit: **2167**

Filing Date: **August 24, 2000**

Examiner: **Mohammad Ali**

For: **PARTIAL MIGRATION OF AN OBJECT TO ANOTHER  
STORAGE LOCATION IN A COMPUTER SYSTEM**

**EXPRESS MAIL LABEL NO: EV670669028US  
DATE OF DEPOSIT: July 1, 2005**

Mail Stop Appeal-Brief Patents  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Sir:

**APPELLANTS' BRIEF PURSUANT TO 37 C.F.R. § 1.192**

This brief is being filed in support of Appellants' appeal from the rejections of claims 1-45 dated 3/2/2005. A Notice of Appeal was filed on 5/2/2005.

**1. REAL PARTY IN INTEREST**

The real party in interest is Microsoft Corporation, recorded as the assignee on 8/24/2000 at reel 011127, frame 0567.

**2. RELATED APPEALS AND INTERFERENCES**

There are no related appeals or interferences

**3. STATUS OF CLAIMS**

Claims 1-45 are pending in the present application, with claims 1, 21, 22, 28 and 30 being the independent claims. Claims 1-45 stand rejected under 35 U.S.C. § 102(e) as

07/06/2005 CCHAU1 00000095 233050 09644667  
01 FC:1402 500.00 DA

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

allegedly anticipated by U.S. Patent No. 5,987,506 (Carter et al.). Claims 1-45 are reproduced in Appendix A, attached hereto, as they stand as of the date of this appeal.

**4. STATUS OF AMENDMENTS**

No Amendments to claims 1-45 have been filed subsequent to the final rejection.

**5. SUMMARY OF CLAIMED SUBJECT MATTER**

Prior to Applicants' invention, the state of the art in hierarchical storage management for files did not cover partial migration of files in most contexts, i.e., it addressed the generic desire of a user to migrate predetermined part(s) of a file from a first location to a second location while retaining other part(s) of the file at the first location. In this regard, the state of the art did not provide a mechanism for specifying those regions of a data stream suited to writes and updates and those regions of a data stream suited to off-line or remote storage. In short, sometimes it is desirable to migrate predetermined part(s) of files to remote storage and to retain other part(s) in local storage and current file servers do not enable native support for and specification of which data to keep and which data to export elsewhere.

The present invention provides such partial migration abilities, and meanwhile, preserves the data relationships of the migrated part(s) to the unmigrated part(s) via file system metadata as a roadmap to reconstruction of the original file. Thus, for example, if part of a Word Processing document were partially migrated to remote storage, for example, the file system of the present invention enables operations on the original document in its entirety without regard to the fact that some of the document may have been migrated to remote storage.

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

**6. ISSUE**

Whether claims 1-45 are unpatentable under 35 USC § 102(e) in view of Carter et al.

**7. GROUPING OF CLAIMS**

**Group I**

Claims 1-21 stand or fall together. Each of these claims calls for the features of “identifying at least one portion of the stream of data for migration to the second storage location; migrating said at least one portion to said second storage location, wherein said migrating includes (A) relocating said at least one portion from the first storage location to the second location and (B) generating additional file system metadata relating to said stream of data; and preserving said stream’s data relationships via said additional file system metadata.”

**Group II**

Claims 22-27 stand or fall together. Each of these claims calls for the features of “an identifier identifying the stream of data for which at least one portion is migrated; data representative of the storage service used in connection with the migration of said at least one portion; and data representative of the memory mappings of said at least one migrated portion.”

**Group III**

Claim 28-29 stand or fall together. Each of these claims calls for the features of an “application programming interface (API) for use in a computer system, whereby a stream of data may register for administration for partial migration techniques according to the method of claim 1.”

**Group IV**

Claims 30-45 stand or fall together. Each of these claims calls for the features of an HSM system that “identifies and migrates at least one portion of said stream of data to a target storage location according to pre-set criteria and generates metadata for the description of data relationships of said at least one migrated portion.”

DOCKET NO.: MSFT-0160 / 142385.1

PATENT

8. ARGUMENT

**Appellant respectfully traverses the Examiner's rejection of claims 1-45 as unpatentable under 35 USC § 102(e) in view of Carter et al. because the claims patentably define over the prior art of record.**

In the Office Action, dated 3/2/2005, the Examiner finally rejected claims 1-45 as unpatentable under 35 USC § 102(e) in view of Carter et al.

Before discussing the specific rejections with respect to each claim, Applicants respectfully submit that the Office Action is improperly combining portions of the description of Carter et al. that have little or insufficient relation to each other with respect to how they are being combined to reject Applicants' claims. According to *In re Arkley*, 455 F.2d 586, 587-88 (CCPA 1972), the "reference must clearly and unequivocally" disclose the claimed subject matter "without any need for picking, choosing, and combining various disclosures not directly related to each other by the teachings of the cited reference." It is important to note that although certain terms such as "data migration," "file system metadata" and "a stream allocation size" may appear in the description of Carter et al., Carter et al. does not describe these notions relating them directly together in any way that would suggest anticipation of Applicants' claims.

For example, Carter et al. speaks of "data migration," "file system metadata," and "a stream allocation size," but does not relate them sufficiently to combine them in the way the Office Action does. First, Carter et al. mentions "data migration" in that the system includes "a migration controller for selectively moving portions of the addressable memory space between the local persistent memory devices of the plural computers." Col. 16, lines 33-35. Carter et al. also mentions "metadata for a directory page," included in a "page header" Col.

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

26, lines 1-2. However, there is no relation made in Carter et al. between this “metadata for a directory page” and the operation of the “migration controller.” Second, Carter et al. mentions that “a stream allocation size reports the total allocation size.” Col. 13, line 14. However, there is no relation made in Carter et al. between this “stream allocation size” and the “metadata for a directory page,” of Carter et al. with respect to data migration. In fact, when mentioning a “stream allocation size reports the total allocation size,” Carter et al. is referring to automatically creating replicas of data streams (Col. 13, lines 10-12), and not data migration.

The final Office Action presented no reply to the argument above presented in Applicants’ previous response.

Thus, Applicants submit that the Office Action is combining portions of Carter et al. in an improper manner to establish anticipation as will become more evident in the discussion of individual claim rejections below.

**Claim 1**

With respect to claim 1, the Office Action alleges Carter et al. discloses “identifying at least one portion of the stream of data for migration to the second storage location” at column 12, lines 3-7. Here, Carter et al. speaks of “desired portions of the data stream,” however, Carter et al. does not specifically teach identifying these portions of a data stream for migration to the second storage location. Carter et al. discloses using these “desired portions of the data stream” generally as part of the “input to the services” to operate at the file level (column 12, lines 3-7).

In response to the argument above, the final Office Action quotes the same section of Carter et al. as above and also adds the following language from Carter et al. “...Further the system can include a migration controller for selectively moving portions of the addressable

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

memory space between the local persistent memory devices of the plural computers. The migration controller can determine and respond to data access patterns, resource demands or any other suitable criteria or heuristic. The migration controller can balance the loads on the network, and move data to nodes from which it is commonly accessed” Col. 16, lines 32-40. Although the language above from Carter et al. mentions a data migration controller, there is no correlation made between the “desired portions of the data stream” that are part of the “input to the services” that a user may utilize to access a data stream (Col. 11, lines 64-67), and the data migration controller “selectively moving portions of the addressable memory space.” These are two different concepts that are not connected in a fashion in which the Office Action purports them to be.

The Office Action also alleges Carter et al. discloses “migrating said at least one portion to said second storage location wherein said migrating includes (A) relocating said at least one portion from the first storage location to the second location and (B) generating additional file system metadata relating to said stream of data” at Col. 3, lines 7-9 and Col. 11, lines 34-35.

With respect to Col. 3, lines 7-9, Carter et al states it is an “object of the invention to provide a globally addressable storage system that employs data migration and replication across interconnected network boundaries and among remote access computers.” Although this is a general description of an object to employ data migration, no mention is made of “generating additional file system metadata” relating to any particular stream of data.

With respect to Col. 11, lines 34-35 of Carter et al., “file system metadata” is used in Carter et al. “to describe a file” and to “represent the file in the file system” (Col. 11, line 38). Although this metadata is associated with a file to identify it, Applicants respectfully submit



**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

that Carter et al. does not teach generating any “*additional* file system metadata” relating to said stream of data.

In response to the argument above, the final Office Action submits that this “*additional* file system metadata relating to said stream of data” is taught by Carter et al. as “...each directory page 120 includes a page header 322 that includes attribute information for that page header, which is typically metadata for the directory page, and further includes directory entries such as the depicted directory entries, 324 and 326, which provide an index into a portion of the shared address space wherein that portion can be one or more pages, including all the pages of the distributed shared memory space.” (Col. 25, lines 66 to Col. 26, line 8).

First, the metadata described above is a general description of attribute information for the directory structure itself, and is not related in Carter to the “portion of the stream of data” purported by the Office Action to be identified for migration by Carter et al. This is evident in looking to the text quoted above by the Office Action, wherein the metadata is described as “metadata for the directory page.” Furthermore, looking to the text previous to the passage above in Carter et al., it is evident that the passage above is referring to Figs. 9 and 10, which “illustrate one example of a directory structure.” Col. 25, line 50. There is no description in Carter of partial migration of this directory structure.

Second, Carter et al. does not teach or suggest that the metadata described above is “additional” file system metadata generated in conjunction with “migrating said at least one portion to said second storage location.” The description of the “metadata for the directory page” described above is provided as a description of the directory structure provided in Carter et al. in order to provide basic access to the shared memory space of Carter et al. (Col. 25, lines 44-51) and there is nothing in Carter et al. that would suggest this file system

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

metadata is “additional” to that present in the system for normal access to the shared memory space.

The Office Action submits that in column 11, lines 60-66 and Fig. 2 et seq, Carter et al. discloses “preserving said stream’s data relationships via said additional file system metadata, whereby said entire stream of data remains accessible to a user of the file system as if said at least one portion of the stream of data were not migrated according to said migrating.” As described above, Carter et al. does not teach generating any “*additional* file system metadata” relating to said stream of data. Here, Carter et al. discusses how data streams are accessed in the system of Carter et al. wherein “Each data stream is individually named so that the user can create or open access to a specific data stream.” Thus, Carter et al. is describing here how a file or data stream is accessed and does not teach “preserving said stream’s data relationships” via any said “additional file system metadata.”

In response to the argument above, the Office Action submits that Carter et al. teaches the above limitation by relying on the following passage from Carter et al. “The stream allocation size still reports the total allocation size in pages required for one replica. The pages backing temporary files, however, are not replicated.” Col. 13, lines 14-15. The Office Action then goes on to state “So there is a relation to preserve the stream’s data with the size.” (see 13, lines 14-15, Carter). However, there is no relation made in Carter et al. between this “stream allocation size” and the “metadata for a directory page,” of Carter et al. with respect to data migration. In fact, when mentioning a “stream allocation size reports the total allocation size,” Carter et al. is referring to automatically creating replicas of data streams (Col. 13, lines 10-12), and not data migration. Therefore, Carter et al. does not teach or suggest “preserving said stream’s data relationships via said additional file system metadata.”

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

Therefore, since all the limitations of claim 1 are not taught by Carter et al., Applicants respectfully submit claim 1 patently defines over this reference. Reversal of the rejection of claim 1 under 35 U.S.C. § 102(e) is thus earnestly solicited.

***Claims 2-20***

Claims 2-20 depend directly from claim 1 and are believed to be allowable for the same reasons. Reversal of the rejections of claims 2-20 under 35 U.S.C. § 102(e) is thus earnestly solicited.

***Claim 21***

The Office Action cites the same arguments for rejection as in claim 1. Thus, with respect to those arguments, Applicants submit claim 21 is allowable for the same reasons given above by Applicants as for claim 1. Reversal of the rejections of claims 21 under 35 U.S.C. § 102(e) is thus earnestly solicited.

***Claims 22-27***

To the extent elements of claim 22 are related to those of claim 1, Applicants submit claim 22 is allowable for the same reasons given above by Applicants as for claim 1.

Additionally, the final Office Action alleges Carter et al. discloses “an identifier identifying the stream of data for which at least one portion is migrated” at Col. 12, lines 3-7. Col. 12, lines 3-7 of Carter et al. state “The file system exports a set of services to operate at the file level. The input to the services are the file object handle anode) or the data stream object handle, and the operation specific parameters, including desired portions of the datastream in byte positions.” Carter et al. speaks of “desired portions of the data stream,” however, Carter et al. does not specifically teach identifying these portions of a data stream for migration to the second storage location. Carter et al. discloses using these “desired

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

portions of the data stream” generally as part of the “input to the services” to operate at the file level (column 12, lines 3-7).

Also, the final Office Action alleges Carter et al. discloses “data representative of the memory mappings of said at least one migrated portion” at Col. 11 lines 60-66, Fig. 2 et seq. However, Col. 11 lines 60-66, Fig. 2 et seq, describe generally how data streams are accessed and named so that “the user can create or open access to a specific data stream.” There is no mention made of “memory mappings” of “said at least one migrated portion.” Therefore, Carter et al. does disclose “data representative of the memory mappings of said at least one migrated portion.”

Claims 23-27 depend directly from claim 22 and are believed to be allowable for the same reasons. Reversal of the rejections of claims 22-27 under 35 U.S.C. § 102(e) is thus earnestly solicited.

***Claims 28-29***

To the extent elements of claim 28 are related to those of claim 1, Applicants submit claim 28 is allowable for the same reasons given above by Applicants as for claim 1.

Additionally, the Office Action alleges Carter et al. discloses “An application programming interface (API) for use in a computer system, whereby a stream of data may register for administration for partial migration techniques according to the method of claim 1.” The data control program of Carter et al. to which the Office Action cites “operates similarly to an interface between a conventional data storage program and a local memory device.” Col. 7, lines 45-47. Carter et al. does not teach using the data control program “whereby a stream of data may register for administration for partial migration techniques.” The data control programs described “can stream data to, and collect data from the

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

addressable shared memory space,” (Col. 7, lines 48-51) but Carter et al. does not describe any *registration* for partial migration techniques.

In response to the argument above, the final Office Action submits that Carter et al. describes exporting “a set of fileset level operations that allow an administrator to manage the filesets through the following type of actions such fileset creation, fileset deletion,...etc (see Col. 10, lines 3-5, Carter).” Carter et al. provides an explanation of what a fileset is by stating “A fileset can be thought of as logically equivalent to a traditional file system partition. It is a collection of files organized hierarchically as a directory tree structure rooted in a root directory tree.” Col. 9, lines 4-6. Thus, according to the definition of fileset provided by Carter et al., the passage of Carter et al. relied upon by the Office Action is providing a description of the fileset level operations on a collection of files (i.e. operations provided on a logical partition of a storage device, or the entire logical volume), as opposed to “a portion of a stream of data.” Furthermore, the specific fileset operations described in Carter et al. are: fileset creation, fileset deletion, fileset enumeration, fileset control and mount export control (Col. 10, lines 6-33), none of which are described as providing registration for a data stream or registration of a data stream for partial migration.

Therefore, Carter et al, does not teach or suggest “An application programming interface (API) for use in a computer system, whereby a stream of data may register for administration for partial migration techniques according to the method of claim 1.”

Claim 29 depends directly from claim 28 and is believed to be allowable for the same reasons. Reversal of the rejections of claims 28-29 under 35 U.S.C. § 102(e) is thus earnestly solicited.

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

***Claims 30-45***

To the extent elements of claim 30 are related to those of claim 1, Applicants submit claim 30 is allowable for the same reasons given above by Applicants as for claim 1.

In addition, citing column 4, line 9-14 of Carter et al., the Office Action alleges Carter et al. teaches an HSM that “generates metadata for the description of data relationships of said at least one migrated portion” Col. 9, lines 9-14 state in part that the “...globally addressable data storage system preferably replicates and/or migrates data stored in the persistent data storage device among two or more of the computers.” However, there is no mention made of any particular migrated portion, nor any metadata generated in conjunction with any migrated portion.

Therefore, Carter et al. does not teach or suggest generating such metadata in Carter et al. “for the description of data relationships of said at least one migrated portion.”

Claims 31-45 depend from claim 30 and are believed to be allowable for the same reasons. Reversal of the rejections of claims 30-45 under 35 U.S.C. § 102(e) is thus earnestly solicited.

Therefore, since all the limitations of claims 1-45 are not taught by Carter et al., Applicants respectfully submit claims 1-45 patentably define over this reference.

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

***Conclusion***

Applicants thus submit that claims 1-45 patentably define over Carter et al. For all the foregoing reasons, Appellant respectfully requests that the Board reverse the rejection of claims 1-45.

Respectfully submitted,



Date: July 1, 2005

---

Jeremiah J. Baunach  
Registration No. 44,527

Woodcock Washburn LLP  
One Liberty Place - 46th Floor  
Philadelphia PA 19103  
Telephone: (215) 568-3100  
Facsimile: (215) 568-3439

DOCKET NO.: MSFT-0160 / 142385.1

PATENT

**APPENDIX A**

***Claims on Appeal***

1. A method for migrating a portion of a stream of data having associated file system metadata via a file system from a first storage location to a second storage location in a computer system, comprising:

identifying at least one portion of the stream of data for migration to the second storage location;

migrating said at least one portion to said second storage location, wherein said migrating includes (A) relocating said at least one portion from the first storage location to the second location and (B) generating additional file system metadata relating to said stream of data; and

preserving said stream's data relationships via said additional file system metadata, whereby said entire stream of data remains accessible to a user of the file system as if said at least one portion of the stream of data were not migrated according to said migrating.

2. A method according to claim 1, wherein said first storage location and said second storage location are located on different volumes.

3. A method according to claim 1, wherein said identifying of said at least one portion for migration includes identifying said at least one portion according to pre-set criteria.

4. A method according to claim 3, wherein said identifying of said at least one portion for migration includes specifying the size of an archive unit.

5. A method according to claim 3, wherein said identifying of said at least one portion for migration includes specifying the size of a region of updates.

6. A method according to claim 3, wherein said identifying of said at least one portion for migration includes specifying a memory allocation limit for the stream of data applicable to said first storage location.

7. A method according to claim 6, wherein said moving of said at least one portion is performed without exceeding said memory allocation limit.



**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

8. A method according to claim 1, wherein the stream of data has at least one identifiable region of updates.

9. A method according to claim 1, wherein said identifying of said at least one portion for migration includes identifying a type of stream of data.

10. A method according to claim 9, wherein said type of stream of data is an append-only file.

11. A method according to claim 9, wherein said type of stream of data is a first storage block write only file.

12. A method according to claim 1, wherein said second storage location is a sequential access medium (SAM).

13. A method according to claim 1, wherein said first storage location is a local location and said second storage location is a remote location.

14. A method according to claim 1, wherein said first storage location is a non-secure storage location and said second storage location is a secure storage location.

15. A method according to claim 1, wherein said first storage location is an on-line location and said second storage location is an off-line location.

16. A method according to claim 1, wherein said preserving the data relationships of said stream includes generating metadata for description of said relationships.

17. A method according to claim 16, wherein said metadata for description of said relationships is formatted according to a scriptable interface capable of being incorporated into World Wide Web components.

18. A method according to claim 16, wherein said metadata for description of said relationships is formatted according to at least one of extensible markup language (XML), distributed component object model (DCOM) and Java.

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

19. A method according to claim 1, wherein the storage for said at least one portion in said first storage location is freed for use by the system after said at least one portion is moved to said second storage location.

20. A method according to claim 1, wherein said stream of data is a sparse file.

21. A computer-readable medium having computer-executable instructions for instructing a computer to perform the method recited in claim 1.

22. A data structure stored on a computer-readable medium for storing metadata relating to migration characteristics of a stream of data wherein at least one portion is migrated via a file system from a first storage location to a second storage location wherein said migration includes relocation of the at least one portion from the first storage location to the second location and generation of additional file system metadata relating to the stream of data, comprising:

an identifier identifying the stream of data for which at least one portion is migrated;  
data representative of the storage service used in connection with the migration of said at least one portion; and

data representative of the memory mappings of said at least one migrated portion,  
whereby said entire stream of data remains accessible to a user of the file system as if said at least one portion of the stream of data were not migrated.

23. A data structure stored on a computer-readable medium according to claim 22, further comprising temporal data relating to a time of migration of said at least one portion of said stream of data.

24. A data structure stored on a computer-readable medium according to claim 22, wherein said data structure is stored according to the format of a scriptable interface capable of being incorporated into World Wide Web components.

25. A data structure stored on a computer-readable medium according to claim 22, wherein said data structure is stored in at least one of extensible markup language (XML), distributed component object model (DCOM) and Java formats.

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

26. A data structure stored on a computer-readable medium according to claim 22, wherein said data structure is stored in a jet database.

27. A modulated data signal for carrying information that encodes a data structure as recited in claim 22.

28. An application programming interface (API) for use in a computer system, whereby a stream of data may register for administration for partial migration techniques according to the method of claim 1.

29. An API according to claim 28, whereby said interface provides a common way to generate and store metadata in connection with the partial migration of streams of data to secondary storage.

30. A computer system including a file system for migrating a portion of a stream of data having associated file system metadata from a first storage location to a second storage location in a computer system, wherein said migrating means to relocate the at least one portion from the first storage location to the second location and generate additional file system metadata relating to said stream of data, comprising:

a hierarchical storage management (HSM) system for administering a stream of data for partial migration; and

a source storage location having a stream of data stored thereon being serviced by said HSM system;

wherein said HSM system identifies and migrates at least one portion of said stream of data to a target storage location according to pre-set criteria and generates metadata for the description of data relationships of said at least one migrated portion, whereby said entire stream of data remains accessible to a user of the file system as if said at least one portion of the stream of data were not migrated according to said migrating.

31. A computer system according to claim 30, wherein the HSM system specifies the size of an archive unit.

32. A computer system according to claim 30, wherein the HSM system specifies the size of a region of updates.

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

33. A computer system according to claim 30, wherein the HSM system specifies a memory allocation limit for the stream of data applicable to said source storage location.

34. A computer system according to claim 33, wherein the HSM system moves at least one portion of the stream of data such that said memory allocation limit is not exceeded.

35. A computer system according to claim 30, wherein the HSM system identifies a stream of data that has at least one identifiable region of updates.

36. A computer system according to claim 30, wherein said identifying by said HSM system of said at least one portion includes identifying a type of stream of data.

37. A computer system according to claim 36, wherein said type of stream of data is an append-only file.

38. A computer system according to claim 36, wherein said type of stream of data is a first storage block write only file.

39. A computer system according to claim 30, wherein said target storage location is a sequential access medium (SAM).

40. A computer system according to claim 30, wherein said source storage location is an on-line location and said target storage location is an off-line location.

41. A computer system according to claim 30, wherein said source storage location is a non-secure storage location and said target storage location is a secure storage location.

42. A computer system according to claim 30, wherein said metadata is formatted according to a scriptable interface capable of being incorporated into World Wide Web components.

43. A computer system according to claim 30, wherein said metadata for description of said relationships is formatted according to at least one of extensible markup language (XML), distributed component object model (DCOM) and Java.

**DOCKET NO.:** MSFT-0160 / 142385.1

**PATENT**

44. A computer system according to claim 30, wherein the HSM system frees for use the storage for said at least one portion in said source storage location after said at least one portion is moved to said target storage location.

45. A computer system according to claim 30, wherein said stream of data is a sparse file.